PC Support Advisor

The Essential Resource for PC Support Professionals

Update 71

Understanding Disk Compression
Cheaper than a new hard disk, but riskier?....3
On-the-fly disk compression is good news for users, but not always such a good deal for those who have to support it. Leigh Edwards explains the benefits and risks involved.
File: A0818.1

Understanding UNIX Commands
Need to know your cd from your ls?....9
If you know your way round MS-DOS but you now need to perform basic operations on a UNIX machine, Julian Moss explains the differences.
File: O0819.1

Hire, Buy Or Lease: How To Decide
It's not always simple to decide....15
With technology changing so fast, and with today's software requiring more and more powerful machines, is outright purchase of PCs always the best way to go? Tim Phillips offers some advice.
File: H0213.1

Supporting OS/2
It's DOS and Windows, but not as we know it....19
Tim Sipples presents a selection of the most frequently asked questions about IBM OS/2.
File: O0720.1

On This Month's Utility Disk

WINZIP - Windows front end to PKZIP

Utilities to unpack tar and .Z files under MS-DOS

"I hate Windows" - says Robert Schifreen.

Until a few months ago, I'd never been a serious Windows user. I'd spent over a decade with MS-DOS, right from version 1.1 upwards, so I was quite at home with the command line. WordPerfect 5.1 for DOS did everything I needed.

Thus, the PC on which I did almost all my work, including the preparation of PCSA, contained a 40 MB hard disk and a whopping 1 MB of RAM.

My backup machine also had a 40 MB hard disk. Its processor was slower than the main machine, but its 8 MB of RAM meant that I could load Windows to check that the programs about to go onto the PCSA Utility Disk worked properly and were free of viruses. The backup PC also had a CD-ROM drive, so I had the luxury of being able to install Windows twice a week, or whenever I wanted to get back to a clean installation.

Some six months ago, my life changed. I was asked to develop some Windows help files, which meant that my main machine would have to switch from DOS to Windows. Out went the 1 MB, to be replaced by 16. Alongside the 40 MB drive I installed a shiny new 420 MB IDE device.

In the old days, I only had room for one Windows app on my machine. Now there was plenty of space so out came Excel, Word for Windows, WordPerfect for Windows, Lotus Organiser, the help file compilers, and anything else I could lay my hands on. After a few days, it was time to do a backup. I was, after all, responsible for almost half a billion bytes, many of which I just couldn't afford to lose.

Now, one of the nice things about running a DOS machine is that you know where everything is. WordPerfect sits in the WP51 directory, and it's easy to put all your document files in a separate place. Indeed, I had separate document directories for each issue of PCSA. So backing up was easy - just point Norton Backup at my C:\DATA directory tree and say "back up everything in there, please". No need to waste time backing up the program files - they're easy to re-install from floppies.

With Windows, of course, it's different. Windows applications like to spread themselves all over the machine. They alter the configuration files. They put their DLLs in \WINDOWS and their fonts in \WINDOWS\SYSTEM. Some of them put their .INI file in \WINDOWS, while some add their data to your own WIN.INI. The main application files also tend to get installed deep in the \WINDOWS directory tree, rather than hanging off the root.

After a couple of months, my WINDOWS and SYSTEM directories had increased in size by 300%.

This whole situation is, of course, total madness and Microsoft should be ashamed of itself for allowing it to happen. Backing up an application becomes impossible, because no one knows where the components of the application are located. Removing an app is just as tricky, and if you delete a DLL that is also used by another application you end up with a program that won't run.

Thankfully, Microsoft has finally admitted that it could have done things better. Windows 4.0 will keep its program directories for its own use, and applications (including DLLs and fonts) will, it appears, go elsewhere. It's one very good reason to consider upgrading your users as soon as version 4.0 ships.

PCSA

PC Support *Advisor*

PC Support *Advisor* is a publication of **International Technology Publishing**.

ISSN: 1031-3966

Editor: Robert Schifreen

Editor in Chief: Les Bell

Publisher: Chris Carvan

Email:
pcsa-ed@oakworth.demon.co.uk

Copyright:
All material is *Copyright © 1994 International Technology Publishing*, and must not be reproduced in part or full, by any means, without the written permission of the publisher.

Liability:
The information, opinions, interpretations and directions herein are based on the best information available, but their completeness and accuracy cannot be guaranteed.

Responsibility:
Responsibility for the views and opinions expressed in this publication lies with the publisher, International Technology Publishing, and not with any agents of the publisher nor the editors or contributing editors.

International
Technology
Publishing Inc.
North American Subscriptions
Suite 200
555 De Haro Street
San Fransisco, CA94107
USA

Telephone: (415) 255 1295
Facsimile: (415) 255 8496

International
Technology
Publishing Inc.
European Subscriptions
13 Callcott Street
London
W8 7SU
United Kingdom

Telephone: (071) 724 9306
Facsimile: (071) 706 3296

International
Technology
Publishing Pty Ltd.
SE Asian Subscriptions
Level 36 Hong Leong Blding
16 Raffles Quay
Singapore
0104

Telephone: (65) 321 8984
Facsimile: (65) 322 8558

International
Technology
Publishing Pty Ltd.
Australasian Subscriptions
Suite 511
185 Elizabeth Street
Sydney, NSW, 2000
Australia

Telephone: (02) 261 4683
Facsimile:(02) 261 4741

Understanding Disk Compression

On-the-fly disk compression is good news for users, but not always such a good deal for those who have to support it. Leigh Edwards explains the benefits and risks involved.

The cost per megabyte of hard disk space is still dropping steadily, but today's obese software has increased the pressure on available disk space and it sometimes seems we're fighting a losing battle.

Decisions to upgrade or replace systems are inevitable in most organisations and the pressure to act usually comes from users who are hitting system limits. It can be expensive, or even impossible, to physically increase disk capacity in some systems.

Disk compression software provides an inexpensive alternative. However, the recent courtroom dramas and reports of reliability problems with DoubleSpace and others have introduced doubts about the wisdom of such a decision. So is it now safe to take the plunge? If you do, what products are there left to choose from and what are the associated support issues? Hopefully, this article will give you the answers.

Note, by the way, that I'm talking about disk compression, rather than file compression as performed by PKZIP, LHARC and others.

History

Disk compression first became available for PCs in the mid 1980s. One of the first products was FreeSpace, from Ashton Tate. It was aimed at dBASE users, but never caught on in a big way. There really was no major need for such technology at the time.

In 1991, Digital Research's DR-DOS 6.0 was bundled with SuperSTOR, thus becoming the first version of DOS to include compression technology as standard (though installation was optional). Other real-time compression software was available, for example, Stacker from Stac Electronics; DoubleDensity from Abacus Software; DoubleDisk Gold from Vertisoft Systems; SuperSTOR from AddStor and XtraDrive from Integrated Information Technology (IIT), but it was not widely adopted.

However, with the introduction of DoubleSpace, the free compression software integrated into, and bundled with, MS-DOS 6.0, many more people were tempted into the compression arena.

Almost immediately, Microsoft was inundated with complaints and reports of bugs. While some used the software without apparent incident, many others quickly encountered problems. Microsoft was quick to vigorously defend DoubleSpace and DOS 6.0, laying the blame at the doors of programmers who had written badly behaved applications.

Eventually, Microsoft released MS-DOS 6.2 in November 93. Officially, this was simply a maintenance release of DOS and was certainly not a bug-fix for DoubleSpace. But users and the press were not slow to notice that there was little new in 6.2 apart from a new version of the DoubleSpace driver plus extra utilities designed to protect data on DoubleSpaced drives. One also wonders why DOS 6.2 was free of charge to users, when all updates had previously been charged for.

All the attendant publicity served to create anxiety among current and potential users about the wisdom of adopting this type of software. It has also stimulated renewed efforts on the part of manufacturers to build in additional safety features, explore new approaches to compression and to make the technology even more invisible to the user.

Early in 1993 Stac won its case against a competitor, IIT, who

> *"Using compression isn't quite as good as a new hard drive, but it can be very inexpensive when compared to adding or upgrading physical drives. With some systems, in particular notebook PCs, compression may be the only viable or economic option."*

conceded infringement of Stac patents. Stac then engaged in a similar battle with Microsoft and were awarded $120 million in damages. In a real life David and Goliath ending, the federal jury decided that Microsoft had infringed two of Stac's data compression patents. However, Microsoft won a counter suit and was awarded $13 million against Stac Electronics for reverse engineering MS-DOS to discover how to implement the undocumented DBLSPACE.BIN loader mechanism, which Microsoft regarded as a trade secret.

As an interim step after the legal decision, Microsoft released MS-DOS 6.21 with DoubleSpace removed. In June 1994, MS-DOS 6.22 was announced, with DoubleSpace replaced by a new program called DriveSpace.

Then, in the same month, Microsoft and Stac settled their differences. Microsoft bought 15% of Stac and gained the rights to use all current and future Stac technology in all current and future OS products. And all for less than the $120m fine (which they don't now have to pay).

If you buy MS-DOS today, you will receive either DriveSpace (6.22) or DoubleSpace (6.20), depending on where you obtain your copy from. Microsoft is claiming full compatibility between these two compression products.

How It Works

File compression software has been around for many years in the form of utilities such as PKZip, ARC, LHARC and so on. Utilities like these have proven to be invaluable for crunching up sets of files or squeezing software onto a minimum number of floppies. But, since they do not work in the background, or in conjunction with application software, they are not the sort of thing for the average PC user. PKZIP's command line is not exactly friendly, and although there are Windows front ends (such as the excellent WinZip shareware package) it's still not ideal. Enter Stacker, SuperSTOR, DoubleSpace, XtraDrive etc which use fundamentally the same data compression algorithms as the manually operated software. Incorporated at a low level into the operating system, they work in the background in a transparent manner to create the illusion of a conventional filing system with about twice its original capacity.

> *"On page 11 of the Stacker 4.0 manual, step 3 in the command sequence for creating an emergency start up disk is wrong. Use CONFIG A: instead of the command syntax given."*

Algorithms

On-the-fly data compression systems use a variety of lossless algorithms, the most common of which are the dictionary-based Lempel-Ziv (LZ) and the patented Lempel-Ziv-Welch (LZW) variant. Stacker calls their patented system LZS for LZ Stacker. LZ type algorithms are generally quick at compression and decompression, which is critical for non intrusive operation.

Lossless algorithms are those in which absolutely nothing is discarded during the compression or expansion processes, and can achieve, on average, compression of around 2:1. Lossy compression, on the other hand, is where some degree of data loss may be acceptable, eg, JPEG used in the compression of images. Lossy compression systems can frequently reduce a file to less than 1% of its size, using special techniques (such as fractal maths) to decompress the file into an acceptable approximation of what it looked like before compression. Although this is fine for graphics, it's obviously not suitable for use with other data.

The LZ schemes use symbols to represent recurring strings in the original data and a dictionary maps the symbols to the strings they represent. Decompression tends to be faster than compression, and since disk reads are far more common than writes, this works in favour of good performance.

The degree of crunching achievable depends upon the type of input data since the method relies upon removing redundancy. Executable files or pre-compressed archive files such as ZIP, LZH, ARC or GIF will not compress much at all, but FAX, bitmaps or sparsely populated database files will be reduced considerably. Text and spreadsheet files usually fall somewhere in between. An average of 2:1 compression is usual but it depends upon the distribution of file types on the disk.

Stac now claims to have exceeded the 2:1 barrier and improved it to an average of around 2.5:1. This has been achieved through special handling of disk sector boundaries, which are described below.

Calculations

It is very important to remember, when looking at any free space figures reported on a compressed drive, that they are only estimated figures based upon the average ratio for the rest of the drive. For example, there is little point (other than for neatness) in storing .ZIPped archives on a compressed volume. You don't gain twice. This apparent paradox takes some careful thought the first time through. With 1 MB of actual space left on a disk, a DIR or CHKDSK command might report upwards of 1.8 MB free. If you try to copy a 1.5 MB text file onto the disk it will probably fit, because text files often shrink by half or more during compression, but a 1.5 MB .ZIP file will almost certainly not fit because it is already compressed and will not shrink any more.

Disk Compression

This is confusing for humans, but often even more confusing to applications. When an application needs to save a data file, it checks to see whether it will fit. Things get very difficult when the disk starts lying to the application about how much space is actually available, and this can lead to all sorts of problems.

Drivers

All of the systems use drivers to trap disk sector reads and writes, and decompress or compress the data coming from or going to the disk. Typically they (except for XtraDrive) hold all the data from a compressed drive in a single huge DOS file (Microsoft calls it a CVF, or Compressed Volume File) and then create a new logical drive for the remaining uncompressed data. A driver swaps the drive letters to avoid having to adjust all the paths. So, in a system with one hard drive, compressing drive C: may result in a new logical drive D: which contains some data which may need to remain outside of the compressed space, eg, DOS system files, CONFIG.SYS, AUTOEXEC.BAT, Windows permanent swap file etc. Drive C:, containing the majority of the files, can now be accessed as normal. (Many compression systems won't allow the DOS system files, or AUTOEXEC.BAT or CONFIG.SYS, to be stored on a compressed drive. This is because the driver to handle the compressed drive won't be loaded until DOS has finished loading.)

For And Against

Using compression isn't quite as good as a new hard drive, but it can be very inexpensive compared to adding or upgrading physical drives. With some systems, in particular notebook PCs, compression may be the only viable or economic option and portable systems are where the capacity limits can hit hardest.

Reduced disk performance may be an issue with software compression, depending on the system specification, although you may get some speed improvements under certain circumstances. With compression in use, fewer sectors have to be read off the disk for a given amount of data. In the early days additional specialised hardware was required to get acceptable speed, but with 386 systems and upwards this is no longer necessary. However things may be turning full circle again - see "The Future" below.

Problems

Data corruption may be more frequent and recovery may be more difficult or time consuming. Often this is due to other software clashing with the cache memory or compression drivers. All the vendors are endeavouring to improve this aspect of their products and the system utility manufacturers are quick to offer compatible tools for checking and fixing compressed drives.

How They Differ

As outlined earlier, a high degree of transparency is an essential feature of this type of utility. The necessary device handlers should not rob the system of too much memory and shouldn't cause awkward changes to the system configuration. Earlier implementations of compression systems created some synchronisation problems with the AUTOEXEC.BAT and CONFIG.SYS files and the potential for DOS memory shortages. Current compression systems are all able, using a variety of approaches, to pre-load their driver software in the very early stages of system boot, ie, before CONFIG.SYS is processed, and are able relocate the drivers to upper memory.

XtraDrive alters the boot track of a hard disk and routes disk function calls through its own device driver and interrupt handler at the BIOS level. In fact the disk driver is loaded before DOS and makes the system even more invisible and easier to use. However, loss or damage to the XtraDrive driver could prevent access to the entire hard disk. Despite this, IIT make great promises for the safety of their product, even claiming that no data loss would occur should the installation process be interrupted, though this is not the sort of thing I would readily contemplate putting to the test.

XtraDrive keeps two copies of a compressed drive's FAT, which are compared on boot up. If a mismatch is found a utility called VMU - Volume Maintenance Utility - works out which version is correct and attempts to carry out repairs.

With MS-DOS 6.0, Microsoft introduced DoubleSpace with a new pre-loading mechanism, via a new version of IO.SYS, to load a device driver called DBLSPACE.BIN that mounts any compressed drives it finds specified in a DBLSPACE.INI configuration file. During the later processing of the CONFIG.SYS file, DBLSPACE.SYS executes, which can relocate the DoubleSpace driver to upper memory.

Starting with version 3.1, and continuing on with 4.0, Stacker uses the same mechanism - this was the basis of Microsoft's counter claim against Stac Electronics. Even though Stac claimed to have written its own DBLSPACE.BIN driver from scratch, Microsoft was successful in claiming reverse engineering of MS-DOS and subsequent infringement of one of their trade secrets. Microsoft however, failed in its attempt to permanently ban Stac from continuing to use the MS-DOS pre-load mechanism.

SuperSTOR/DS, included with IBM's PC-DOS 6.1, uses the same preload mechanism, as does Novell DOS 7 (previously DR-DOS) which now includes Stacker. Before being taken over by Novell, Digital Research's DR-DOS included a version of SuperSTOR that used yet another loading variant via a DCONFIG.SYS file. The AddStor supplied version of SuperSTOR Pro, which soon will no longer be available, did not pre-load. (AddStor recently went out of business. IBM is now bundling Stacker with PC-DOS. If you still run SuperSTOR on any PCs, you will no longer be able to get support in case of trouble so it's a good idea to remove it.)

While Stacker, DoubleSpace and SuperSTOR hold all the compressed data in one large file, which could mean loss of all data if it becomes seriously damaged, XtraDrive takes a very different approach. Each file is

> *"Both DoubleSpace and Stacker take a considerable amount of time to install. On a 25 MHz 486DX system, with a caching disk controller, Stacker 4.0 took a total of three hours."*

completely separate from all the others. If one file becomes damaged, the others remain intact. XtraDrive creates two drives from one, called the primary and secondary drives. The primary drive contains all the original data and the secondary drive is given a new logical drive letter and contains the additional space. XtraDrive is also able to compress the DOS system files and the Windows swap file.

DoubleSpace includes the MRCI - Microsoft Real-Time Compression Interface. This allows an application to query the operating system for the presence of a compression server. If such a server exists, the application can make use of its compression/decompression engine, eg, a communications package could call the routines before transmitting data across a link. DOS 6.x provides an LZ based compression server which is the engine which DoubleSpace uses. The new MS Backup and MS Restore utilities provided with DOS 6.x can make use of the DoubleSpace compression engine, thus using fewer floppies in a backup set.

MS-DOS 6.2 includes DoubleGuard, which is loaded automatically by DoubleSpace (although it can be disabled), and gives some extra protection against disk corruption by checking the integrity of some of DoubleSpace's data structures. If DoubleGuard detects that the memory used by DoubleSpace has been corrupted by other errant software, it will halt the system, but unfortunately, without giving you the option of saving your work.

There is an undocumented override to this, which you use at your own risk. Pressing the grey + key on the numeric key pad, on getting a DoubleGuard alarm, will allow you to continue.

ScanDisk is a utility for the detection, diagnosis and repair of disk errors on uncompressed and DoubleSpaced drives. As part of the improved compression safety in MS-DOS 6.2, ScanDisk is run automatically by DoubleSpace before a new compressed drive is created. This feature was added after reports in the press that DoubleSpace failed to ignore bad sectors when creating a CVF.

Write back caching is now off by default in SmartDrive. Previously on by default, this could cause problems if users turned off their machines or pressed Ctrl-Alt-Del while in Windows. Also SmartDrive flushes the cache before passing the reboot request on to the operating system when Ctrl-Alt-Del is pressed (this does not work if you have VSAFE loaded after SmartDrive in the AUTOEXEC.BAT).

Sector Boundaries

Stacker 4.0 is the first disk compression software to ignore normal disk sector boundaries. By handling the management of disk sectors and clusters itself, plus the clever use of file extents, Stacker 4.0 uses disk space more efficiently. Stac calls this system SmartPack. Normally, DOS allocates disk space in whole clusters, which consist of groups of disk sectors. A small text file containing only a line or two of text takes a whole cluster, which may be up to 32 KB on a 1 GB drive (DOS only has a small space in which to store data about cluster contents, so cluster sizes on large disks have to increase). Under Stacker 4.0, fragmented space is more intelligently handled on a compressed disk, and allows spreading a file across the available smaller units of disk space. (The latest release of Novell NetWare also does this.) Also, Stacker 4.0 can create a compressed volume up to 2 GB whereas DoubleSpace is limited to 512 MB.

Using Novell's DOS Protected Mode Services (DPMS), Stacker 4.0 is able to load its drivers into extended memory where programs are protected from each other, hence providing similar safety measures to Microsoft's DoubleGuard. DPMS also has the benefit, with DOS 6.2, of reducing the conventional DOS memory footprint to less than 2 KB. Loaded high, Stacker 4.0 uses less than 24 KB of upper memory, although Stac claim as little as 17 KB is required. If installed on a system with data previously compressed with DoubleSpace or SuperSTOR/DS, Stacker 4.0 will convert the compressed data to its own format.

Stacker With Windows

An impressive and expanded set of Windows tools is provided with Stacker 4.0. All commands are available from the Windows Toolbox. Many equivalent facilities are available from DOS by entering STACKER at the DOS prompt while in the Stacker sub-directory. Facilities are provided to check data integrity and drive status, defragment drives, tune between speed and compression ratio (over a scale of 1 - 10) and to sound alarms at user defined levels, eg, fragmentation level, available disk space etc. Also, an additional Stacker Tools item is added to the Windows File Manager menu.

Additional safety facilities are included in version 4.0, eg, AutoProtect, which checks the health of the data, and AutoSave, an optional data safety feature which periodically saves key system data in an uncompressed state. It's worth noting, however, that AutoSave can interfere with certain time critical operations, eg, comms file downloads, so it should be turned off prior to this type of operation.

Disk Compression

32-Bit File Access

Stacker 4.0 does not currently support the 32-bit file access feature of Windows for Workgroups 3.11. Stac has recognised this as an important problem and are working to address it. In the mean time they recommend that 32-bit file access (32BFA) should be disabled on systems where Stacker 4.0 is running. In practice I have found that Stacker disables 32BFA during the installation and setup of a compressed drive, but leaves it enabled for other uncompressed drives. Stacker 4.0 is fully compatible with 32-bit disk access.

Safety First

As mentioned earlier, all of the systems except XtraDrive use the equivalent of Microsoft's Compressed Volume File (CVF) to hold the compressed data on the uncompressed host drive. There is the apparent danger that, since all of the compressed data is in one big file, disk crashes involving compressed drives can be more serious and difficult to recover from. All of the compression filing systems maintain detailed mapping information, analogous to conventional File Allocation Tables (FATs), as well as cross-checking information and safeguards to maintain data integrity.

Some systems keep multiple versions of their file maps to make error detection and structure rebuilding easier. The key issue though, is not the exact method used to achieve the end result, but how easily checking can be done and what options are available in the event of a fault occurring.

The utility manufacturers such as Symantec, with Norton Utilities, and Central Point, with PC Tools, are attempting to keep pace with the changes in the various products, in particular with the compression market leaders. For example, Norton Utilities v7 will work fine with DoubleSpace and earlier versions of Stacker, but even the latest incarnation, version 8, of Norton Utilities is not fully compatible with Stacker 4.0. Norton 8 will apparently repair the host drive of a Stacker 4 compressed drive, but it does not understand the internal structure of the compressed drive itself. So, until either Stacker or Norton gives the all clear, I wouldn't even attempt to use them with version 4.

If in doubt, the safest bet is to use the check and repair utilities provided with the system itself, eg, ScanDisk, Stacker Toolbox etc.

Experience gained while providing support and during personal use has led me to conclude that compressed data is not as safe as uncompressed data. My early experiments with DOS 6.0 and DoubleSpace were not encouraging. Being slightly paranoid about data loss, I regularly run disk checking utilities to detect and correct lost chains and other disk corruptions - an unfortunate fact of life with DOS. The frequency with which errors occurred rose sharply on implementing DoubleSpace in DOS 6.0. It was the complete loss of a compressed drive that caused a quick withdrawal from DoubleSpace.

After the release of MS-DOS 6.2 I was tempted to try again. Although the frequency of file errors seemed lower, I again withdrew after another unexplained disk loss. Despite these irritating experiences there was no significant, actual data loss involved because of adequate backup measures. On a related topic, be careful with undelete and DoubleSpace in MS-DOS 6.2 - you may find that a file that has apparently been correctly undeleted is in fact incompletely recovered.

None of my Stacker 3.x users has reported problems to date, but very recently however, I suffered the complete loss of a Stacker 4.0 drive and so once again I went in search of the backup tapes with fingers crossed.

If you decide to proceed with compression, backup the whole system first before installation. Set up or make any required changes to the Windows

Compression Tips

- Consider a bigger or an additional hard disk first. Over all, it may be more economical than using compression.
- Always back up the whole system before installation or removal of compression, or any disk maintenance and repair work.
- Clear off all unused files, particularly any files that are already compressed, eg, .LZH, .ARC, .ZIP, .GIF etc.
- Ensure that a disk integrity check is done first, in particular do a disk surface check.
- If you chose DoubleSpace, avoid DOS 6.0 as it has many problems. Use 6.2.
- Read the installation instructions in detail and carefully follow any instructions to disable TSRs. (Stacker 4.0 now does this automatically.)
- After installation, use the disk integrity checking and defrag utilities regularly, preferably at the end of the day because it can be very time consuming.
- Where the data is very important and/or highly volatile, check the backup frequency and consider increasing it.
- Bear in mind the capacity of the backup systems you have available. Extensive use of compression within an organisation may generate a need to upgrade the some or all of the backup units.
- Get the latest information available for the product you choose. Stac runs a very good BBS in the US (telephone +1 619 431 5956). You will be asked for some personal details including your product serial number (although I found this not to be necessary). Download the self extracting archive file XXX999.EXE where XXX stands for the version you are using, eg, DOS999.EXE, OS2999.EXE etc. It contains all the past and present technical support notes. Also available are updated versions of the defrag and uncompress programs, SDEFRAG2.EXE and UNCOMP2.EXE respectively, which were initially problem prone.

Microsoft has a US BBS on +1 206 936 6735. There are also active support forums on CompuServe for both Stac and Microsoft products.

permanent swap file first. If you haven't already done so, setting up all your Windows users with permanent swap files will improve their system performance and might save you a lot of bother later. Swap files, of the permanent kind, usually have to be on an uncompressed drive (apart from XtraDrive). With Stacker 4.0 you can add a permanent swap file on a compressed drive at a later date using the Stacker Windows Toolbox. The swap file will be set up in an uncompressed form and will reduce the amount of compressed space (estimated) available by a factor of around 2.5 times the swap file size.

Manual Error

On page 11 of the Stacker 4.0 manual, step 3 in the command sequence for creating an emergency start up disk is wrong. Use CONFIG A: instead of the command syntax given. For versions of DOS which pre-load compression, ie, MS-DOS 6.x, PC-DOS 6.1 or Novell DOS 7.0 the procedure on page 10 should be used and supplemented by copying the hidden file STACKER2.BIN from the root directory of the uncompressed drive to the root directory of the start-up disk. Including this file will provide additional speed should you need to boot from this disk.

Installation Time

Both DoubleSpace and Stacker take a considerable amount of time to install. On a 25 MHz 486DX system, with a caching disk controller, Stacker 4.0 took a total of three hours to convert a 294 MB logical drive with 30 MB free to give a 733 MB drive with 429 MB free. The process included preparation of the drive which took nearly two hours, followed by 50 minutes to defrag the disk and then 20 minutes to verify the new Stacker drive. DoubleSpace turned in similar figures on the same system.

Removal

Removing DoubleSpace under MS-DOS 6.0 is a pain, though it is now a lot easier with 6.2. Obviously if you needed DoubleSpace then you've probably got more files and data stored than will fit on the disk in its uncompressed state. So that's the first problem you'll have to solve, if you're going to remove compression. Be sure to back up all your data before starting this operation.

When you've freed up some space on the uncompressed drive, you'll need to move the contents of the compressed drive to the uncompressed drive. (Deleting any Windows permanent swap file while you do this job is a good way of releasing a nice big chunk of disk space.)

Under DOS 6.0, use the DBLSPACE /SIZE command to shrink the size of the compressed drive and then move more files over to the uncompressed drive. Repeat this process until all the files have been removed from the compressed drive. Then use the DBLSPACE programme to remove the compressed drive completely.

A simpler, but possibly more risky, procedure would be to backup all the data, remove the compressed drive, then restore the required data to the uncompressed drive.

Having transferred the data and removed the drive, you now need to stop DBLSPACE.BIN loading during system boot. Edit the CONFIG.SYS file to remove any lines that refer to DBLSPACE.SYS and save the file. Change to the root directory and enter DELTREE /Y DBLSPACE.* and then reboot the system. With MS-DOS 6.2 you simply enter the DBLSPACE /UNCOMPRESS command from the DOS prompt and follow the on screen instructions.

A Stacker drive is easily converted back to uncompressed format from the Windows or DOS toolboxes, providing you have the space available. On removing the last compressed drive Stacker will ask if you wish to remove it from the system configuration entirely.

The Future

Compression is definitely here to stay, in one form or another. If doubts about robustness can be removed, then increasing numbers of users will want to reap the economies and benefits.

Chicago (Windows 4.0) will include compression technology in order to maintain compatibility with DOS 6.x. Windows NT Remote Access Server and MS Flash File Server have made use of DoubleSpace technology. Current and future versions of IBM PC-DOS, will incorporate compression technology. Stacker is obviously very determined to stay in the compression business - they have products available for DOS, OS/2 and Macintosh and are still very active in hardware based compression.

When 808x systems ruled and 286 based PCs were the state of the art, some compression systems incorporated ISA based interface boards which handled the compression and decompression via high speed chips. Software-only solutions were too slow and the hardware based solutions, of that time, were limited by the speed of the ISA bus.

Hardware compression is often found today in tape backup systems, where the users are hardly aware of its presence. By integrating the technology directly into the CPU's I/O bus, in the CPU itself or in the form of dedicated co-processors, data compression could become analogous to floating point maths. Considerable performance enhancements could be achieved using such an approach since data could be moved around in compressed form and only expanded when absolutely necessary. A number of companies, including Stac and IBM are actively exploring new hardware based options.

PCSA

The Author

Leigh Edwards is an independent consultant, widely experienced in establishing and managing computer support groups within a major international pharmaceutical company.

Understanding UNIX Commands

If you know your way round MS-DOS but you now need to perform basic operations on a UNIX machine, Julian Moss explains the differences.

UNIX may not be an operating system that you are currently required to support, but a working knowledge of it can still be useful. One day, you may find yourself having to deal with UNIX-based servers on your network. You will also find a familiarity with UNIX commands useful if you are interested in exploring the Internet - you'll need basic directory navigation skills to explore the ftp sites, for example.

UNIX, like DOS, is likely to be with us for years to come. Though you may not yet have encountered it, it will be surprising if you never have to.

DOS Differences

If you're familiar with DOS, you already understand many of the concepts of an operating system. UNIX pre-dates DOS by about 10 years, and was written as a multi-tasking, multi-user operating system for mainframe computers. Many of the ideas behind DOS were borrowed from UNIX, though as a single-user, single-tasking operating system for microcomputers DOS is necessarily far less sophisticated.

> *"If you're familiar with DOS, you already understand many of the concepts of an operating system. UNIX pre-dates DOS by about 10 years, and was written as a multi-tasking, multi-user operating system for mainframes."*

case sensitivity

There are a few obstacles which tend to trip up the DOS user who is starting to learn UNIX. One is case sensitivity. DOS isn't case-sensitive, but UNIX is. So while **`login`** is a valid command and **`/bin`** a directory, **`LOGIN`** and **`/BIN`** probably won't exist and will result in an error message. UNIX passwords are also case sensitive, which is worth remembering if you have an account on a UNIX system and you're wondering why your password appears not to work.

Another common source of trouble is that UNIX uses a forward slash as a directory separator. The backslash is used for other purposes in the UNIX command shell, and so can't be used in pathnames.

As a DOS user you are probably conditioned to using backslash when typing a directory name or a path, so you'll need to make a conscious effort to break the habit.

UNIX commands use the dash character, **`-`**, to denote command line switches. DOS uses the forward slash, although some DOS utilities (such as the compression tool PKZIP) parse their own command lines instead of letting DOS do it for them, and do use a UNIX-like dash.

UNIX command switches generally appear after the command name, and before any other arguments. This is different from DOS, where they often appear at the end of the command line. Note, too, that under UNIX, switches are combined, so if you wanted to invoke both switches a and b in a command you would use **`-ab`** or **`-ba`** rather than **`-a -b`**.

The File System

Like DOS, UNIX stores files in a hierarchical directory system. The directory names are usually chosen by the user or system administrator to reflect the content of the files stored there. However, there is greater standardisation of directory names under UNIX than you find under DOS.

Directory **`/bin`** is used to hold the operating system binary files. Binary files are what UNIX users call executable files. UNIX system binaries do a similar job to the COM and EXE files in the \DOS directory on a PC - the DOS "external commands". Directory **`/etc`** contains miscellaneous UNIX configuration files, while **`/dev`** is used to hold device drivers. User directories are usually located in branches off the **`/usr`** directory.

UNIX file names are not restricted to the 8.3 format of DOS, and may contain pretty well any printable character that doesn't have a special significance to the UNIX command shell. This means that filenames like **`read.this.file.now`** are perfectly legitimate. Such filenames may cause problems if you try to use a file transfer

> *"A common source of trouble is that UNIX uses a forward slash as a directory separator. The backslash is used for other purposes in the UNIX command shell, and so can't be used in pathnames."*

protocol to copy the files to a DOS machine.

Under UNIX, file extensions are not generally used - as they commonly are in DOS - to indicate what sort of data a file contains. Some UNIX filenames do follow a similar convention, ending with a period and one or more letters. For example, a filename ending in .Z shows that it has been compressed using the UNIX "compress" utility, and must be uncompressed (using "uncompress") before it can be used. A filename ending in ".tar" is a tape archive format file, while a ".uu" file is a binary file that has been encoded so that it can be sent via electronic mail. However, these extensions are simply a naming convention, and are treated as part of the filename. If you download a file from an ftp site over the internet, onto a DOS machine, you'll need a DOS version of the necessary uncompression utility before you can access the file. *[You'll find some of these on this month's PCSA Utility Disk - Ed].*

One thing that may strike you on seeing a list of files in a UNIX directory is that, without identifiable executable file extensions, there is no obvious way to tell the executable files from the data. Under UNIX, an attribute in the file's directory entry shows whether it can be executed or not. I'll illustrate this later.

As in DOS, wildcards can be used in most UNIX commands to match multiple filenames. The wildcard `?` stands for any single character, while `*` represents zero or more characters. However, UNIX has some more powerful wildcard options that don't have an equivalent in DOS. Square brackets can be used to list specific characters to be included, so `[bcm]at` would match bat, cat and mat, while `[a-m]*` would match filenames beginning with a through m.

UNIX does not distinguish between a filename and an extension, so the period character does not have any special significance. It is just part of the filename. This has an important implication when using wildcards. The string `*.Z` would match all the filenames ending in `.Z`. Note that the wildcard that matches all files is `*` and not `*.*`.

Getting Started

UNIX systems are usually multi-user. To gain access to them you need to log in using the **`login`** command. You'll probably also be required to

Putting It Into Practice

If you have a UNIX box in your organization then you can use it to try out the UNIX commands described in this article. Ask the system administrator to set up a user account for you to experiment in. If the UNIX system is connected to your PC LAN then you may be able to log in from your PC using Telnet, a terminal emulator that works across an Ethernet connection using the TCP/IP protocol.

If you don't have access to your own UNIX box then there are several alternatives. One is to use a UNIX-like shell and set of utilities that runs on top of DOS and enables you to use UNIX commands on your DOS-based PC.

Possibly the closest emulation of UNIX under DOS is provided by Mortice Kern Systems' MKS Toolkit. This can be installed in a variety of configurations. At one extreme, you continue to use COMMAND.COM but can use UNIX commands as an alternative to those of DOS. At the other, you can set up the system to support multiple logins, which give you the option to use the Korn Shell, COMMAND.COM, or Windows as your working shell.

Other possibilities worth investigating are Coherent, from the Mark Williams Company, and Minix, a UNIX clone which is published as a book by Prentice Hall. There are also versions of UNIX in the public domain, including Linux, which comes on a CD-ROM and includes an X-Windows graphical interface. These UNIX clones replace DOS entirely, requiring you to dedicate a system - or at least a partition of the hard disk - for their use.

If you have Internet access from your PC then you can try out some of the commands described here on sites you connect to using Telnet or FTP. However, there will be limitations: you are unlikely to be allowed to create or remove files, for example.

Using FTP you can use **`cd`** to navigate around the directory tree of the system you have connected to, and **`ls`** to display the contents. However, your use of other commands may be restricted.

You can Telnet to many public-access sites, although most have a menu-driven interface so you won't see a raw UNIX prompt. One that doesn't is the UNIX User group archive at src.doc.ic.ac.uk. You can use commands to list files, but you will still not be able to create, modify or delete them.

UNIX Commands

> *"Editing files is one task that's a lot harder to do under UNIX than under DOS. None of the UNIX editors is as easy to use as the full-screen, mouse-driven MS-DOS editor."*

enter a password. If you don't have an account, your system administrator will need to set up one for you.

Once you are logged in, the UNIX system loads the command shell and sets the appropriate terminal emulation (UNIX users may use a variety of types of terminal) along with other preferences such as your home (or start-up) directory. These preferences are set up when a user account is created by the system administrator.

The UNIX command shell is the equivalent of COMMAND.COM. It performs a comparable range of functions: it enables you to execute commands by typing them at the prompt; it lets you specify search paths and use wildcards in filenames; and it supports the use of pipes, redirection and batch files. UNIX batch files are called shell scripts.

One difference between the UNIX shell and COMMAND.COM is that there are no internal commands. Every UNIX command is implemented by running a binary file stored in the directory **`/bin`**. The solution adopted by DOS - to make the most commonly-used commands part of COMMAND.COM - was chosen to make operation more convenient on PCs without a hard disk because it allows the system diskette to be removed from the drive.

On most UNIX systems there is a choice of command shells. The Bourne shell is commonly used for user accounts, but system administrators, programmers and others usually prefer the Korn shell, a superset of the Bourne shell which has more powerful scripting features. (There are DOS versions of these shells, available from companies such as MKS, which are useful if you need to practice your UNIX at your PC. They're also wonderful for UNIX people who need to use a PC occasionally.)

Directory Commands

Once you've logged in, the best way to get a picture of what the system contains is to list some directories. Remember, if you want to experiment and you don't have a UNIX machine handy, almost any ftp site on the internet will do.

Under UNIX, the command to list a directory is **`ls`** or, more fully, **`ls {switches} {path}`**. On its own, **`ls`** lists just the names of the files in the current directory, in multi-column format. This is the equivalent of DIR/W in DOS. The **`ls`** command has many switches, most of which are only needed for special purposes. The switch you will use most often is **`-l`** which produces a long format listing that equates in function, if not in appearance or content, with the standard DIR command.

Figure 1 shows an output from an **`ls -l`** command. The first character on each line shows whether a file is a directory (d) or an ordinary file (-). The characters which follow show the access permissions. These are read (r), write (w) and execute (x). Dashes are used to show where a permission has not been granted.

UNIX supports three classes of user for which separate access permissions may be granted: owner, group and world. There are three groups of rwx, representing the permissions for each of these classes. The owner is the creator of a file; group is a group of users which may be given access to it, and world is everybody else on the system.

Next, the listing shows the number of links to a file. Under UNIX, a file may be stored once on a disk but appear in many directories: each separate entry is one link. When you delete a file it will be removed from the directory, but it will only be deleted from the disk when the number of links to it is zero. The two names which appear next are the owner of the file or directory, and the name of the group that owns the file or directory. The name root often appears here: root is the system user name.

Following the links comes the file size in bytes and the date and time a directory was created or a file was last changed. Finally, the directory or file name is given.

Moving Around

To move around the directory hierarchy you use the **`cd`** command. This works much like its equivalent in DOS: the full format is **`cd {path}`**. If no path is specified, the **`cd`** command returns you to your home directory.

Listing Files

To display files under UNIX you use the **`cat`** command, much as you would use TYPE under DOS. The name comes from the word "concatenate" because, if you specify two or

```
-rw-r--r--  1 root     root        1202 Mar  9 14:08 README
drwxr-xr-x  2 root     root         512 Apr 23  1993 bin
drwxr-xr-x  2 root     other        512 Jun 21 10:46 data
drwxr-xr-x  2 root     root         512 Jun  7 09:32 dev
drwxr-xr-x  2 root     root         512 Aug 24  1992 etc
-rw-r--r--  1 root    root         827 Mar  9 14:48 login.anon
drwxr-xr-x  3 root     root         512 Oct 18  1991 usr
```

Figure 1 - Sample UNIX directory listing

more filenames, `cat` sends the contents of each of the files to standard output, one after the other. The general format is `cat {switches} {filename(s)}`.

For the most part, you can forget about the switches for `cat`: you are only likely to need them in special cases. To list a file to the screen, you simply type: `cat filename`.

If no filenames are specified, `cat` takes its input from standard input - the console. This leads to another common use for `cat`, which is to create simple text files - like using COPY CON: in DOS. If you enter `cat >TEXTFILE` then anything you type, up to the end-of-file character Ctrl-Z, will be directed to a file named TEXTFILE. Note the use of output redirection in this example. The redirection characters in UNIX are the same as in DOS: > for overwriting a file with redirected output, >> to append your redirected output to a file instead of overwriting it and < to redirect input from a file to a command.

To use `cat` to concatenate several files, you would enter `cat file1 file2 file3 >newfile`.

Like the DOS TYPE command, `cat` does not pause a screenful at a time when displaying output. To pause the display, you can pipe the output into another utility: `more`. For example `cat README | more`. (Sound familiar? Guess where the developers of DOS found a lot of their inspiration.)

> *"UNIX may not be an operating system that you are currently required to support, but a working knowledge of it can still be useful. One day you may find yourself having to deal with UNIX-based servers on your network."*

Note that the pipe operator is also identical to that used in DOS. However, the UNIX version of `more` is somewhat more powerful than that under DOS. You can page both forward (by pressing space, F, Ctrl+F or PgDn) and backward (by pressing B, Ctrl+B or PgUp). The reason for the choice of paging keys is that UNIX systems may be accessed from a variety of terminals, some of which may not have a cursor keypad or a Ctrl key.

Using `more` you can scroll forward a line at a time (by pressing Return, j or the down arrow key) and back a line at a time (by pressing k or up arrow). `More` also lets you skip a specified number of lines, set and return to position markers in the text, and even issue operating system commands without leaving the utility. To display a full summary of `more`'s control keys, press the h key which brings up a page of help.

> *"UNIX file names are not restricted to the 8.3 format of DOS, and may contain pretty well any printable character that doesn't have a special significance to the UNIX command shell. Some UNIX filenames may cause problems if you try to use a file transfer protocol to copy the files to a DOS machine."*

Though the example above demonstrates the use of the pipe character, the use of piping is unnecessary because you could more concisely enter `more README` to achieve the same result.

An alternative to `more` on some systems is `pg`. This utility isn't quite as easy to use as `more`: it is driven by entering commands at the : prompt at the foot of the screen. For a list of its options, once again press h.

Examining Files

Often, you simply want a quick look at a file to see what it is. The command `head filename` displays just the first 10 lines of a file before returning to the command prompt. Similarly, the command `tail filename` lists the last 10 lines.

File Maintenance

To move or rename files under UNIX you use the `mv` command. The format is `mv {switches} filename1 filename2`. Alternatively, there's `mv {switches} file1 {file2 file3 ...} directory`.

The first format is used to move or rename a single file. The second format is used to move a list or group of files to another directory. Mostly, you won't need any command switches, but you can include the switch `-i` if you want a prompt before overwriting any file.

There is a third format for the mv command: `mv -r directory1 directory2` moves the entire contents of one directory to another. You can use `-R` in place of `-r`, with identical effect.

Copying Files

The UNIX copy command is `cp`. It has the same formats as the `mv` command: `cp {switches} filename1 filename2` or `cp {switches} file1 {file2 file3 ...} directory` or `cp {switches} directory1 directory2`. Use the `-i` switch if you want a prompt before overwriting anything.

Unlike DOS's COPY, `cp` normally sets the date and time of modification of the target file to the current date and time.

To preserve the original date and time, use the `-m` switch. Use the `-p` switch instead if you also want to preserve the file access modes, owner and group owner.

Recursive Copying

The `-r` switch can be used in the third format of `cp` to copy an entire directory. However, to make an exact clone you should use `-R`, which takes steps to ensure that certain special types of file are duplicated correctly.

To remove a file from a directory, you use the command `rm`. Unlike the DOS DEL command, you can specify several filenames on one command line, so the format is: `rm {switches} file1 {file2 ...}`.

You can include the `-i` switch to obtain a prompt for confirmation before deleting each file. The `-r` or `-R` switch will remove an entire branch of the directory structure if the named files are directories, rather like MS-DOS's DELTREE.

Housekeeping

Under UNIX, users don't need to worry about checking for lost clusters or defragmenting their hard disk. UNIX checks the integrity of the file system when it boots up, before anyone has a chance to log in. If you need to find out how much disk space is used, or how much is free, UNIX provides commands which will do just that.

> *"Once you have logged in, the best way to get a picture of what the system contains is to list the contents of some directories. Remember, if you want to experiment and you don't have a UNIX machine handy, almost any ftp site on the internet will do."*

The command `du` tells you how much disk space is used by a directory and its subdirectories. For example `du /usr/john` would display a list of directories starting at `/usr/john` showing how much space is occupied by each of them. By default, UNIX systems report disk usage in blocks of 512 bytes. To obtain a report in KB, remember to specify the switch `-k` before the path name.

To find out how much disk space is free, use the command `df`. The format of the command is `df {switches} {{device} {filename}}`.

The output from `df` gives two figures: the amount of space free on a device, and the total amount of space both free and used. As with the `du` command, you should use the `-k` switch if you want the space reported in KB instead of blocks.

If no device or filename is given, `df` reports the amount of free space on each of the volumes that is mounted. You can obtain a report of the space available on just one device either by specifying it by name, or by giving the name of a file stored on that device.

To change the access permissions of a file, you must use the `chmod` command. The format is `chmod {switches} mode pathname`. The mode consists of the letters r, w and x (for read, write and execute) preceded by an operator. The - and + operators are identical to those of the DOS ATTRIB command, ie, - removes the access permissions which follow it, + adds them. There is also an = operator, which turns on the access permission specified and turns off all the others.

The operator can optionally be preceded by a letter u, g, o or a, which allows you to specify whether you want to change the user, group or other permissions, or all of them (the default.) A couple of examples should help to make this clearer. To remove the write permission entirely from a file, you would use the command `chmod -w filename`. To enable the group execute permission of a file, you would use `chmod g+x filename`.

Text Editing

Editing files is one task that is a lot harder to do under UNIX than under DOS. UNIX systems usually have a choice of editors available, but none of them is as easy to use as the full-screen, mouse-driven MS-DOS Editor. The standard full-screen editor which you will find on every UNIX system is `vi`. You start `vi` by typing `vi filename`. If `filename` does not exist, `vi` will create it.

`vi` is a modal editor. This means that you are always in one of two modes: add mode, where what you type is entered into the file you are editing, or command mode, where what you type is interpreted as an instruction. To get from add mode to command mode you press the Esc key.

When you start `vi`, you are initially in command mode. If you are trying to create a new file, you can't just start typing. You must start by typing **a**, the add command. This puts you into add mode, and you can then begin to enter text. You may notice the cursor change from an underline in command mode to a block in add mode. This subtle change is the only indication you get

> *"Under UNIX, users don't need to worry about checking for lost clusters or defragmenting their hard disk. UNIX checks the integrity of the file system when it boots up, before anyone has a chance to log in."*

of which mode you are in. If you are ever confused about which mode you are in (and believe me, you will be) just press Esc. This will put you back into command mode. If you are already in command mode, **`vi`** will just beep.

While in add mode, **`vi`** lets you move about using the cursor keys, and generally behaves like a normal full-screen editor. However, there are a few differences which are attributable to the need for **`vi`** to be able to work on basic terminal equipment connected by slow serial lines. For example, if you use the backspace to delete characters, the characters don't disappear from the screen until you go to command mode or move the cursor to another line. This cuts down the time spent refreshing the screen.

In command mode, you can type **`H`** (High) to move the cursor to the first line of the screen, and **`L`** (Low) to move it to the last. These commands will probably be mapped on to the Home and End keys if your terminal has them.

You can scroll down the file by pressing Ctrl+D (Down) or Ctrl+F (Forward). The former moves the text half a screenful at a time, the latter moves a full screen. To go back in the file you can use Ctrl+U (Up) or Ctrl+B (Back). Forward and Back will be mapped on to PgDn and PgUp.

To quickly move the cursor to the start of a line, make sure you are in command mode and then press **`0`**. To move to the end of a line you press **`$`**. You can also skip forward from word to word by pressing **`w`**, and backward by pressing **`b`**.

You can delete text from the cursor position to the start of the next word using the command **`dw`**, and from the cursor position to the end of the line using **`d$`**. You can delete the whole line in which the cursor is placed using the command **`dd`**. Remember, to add text after deleting some you must press **`a`**, to get into add mode, and then begin typing.

To save your work, get into command mode and then type **`:w`**. Note the colon, which is part of the command. If you want to quit **`vi`**, type **`:q`**. You can combine these two commands into one if you wish, by typing **`:wq`**. If you want to quit **`vi`** without saving your changes, you must use the command **`:q!`**.

This subset of **`vi`** commands should be enough to get you by if you need to perform simple tasks like editing configuration files.

Getting help

If you get stuck during a UNIX session, and can't remember the format of a command or which switches to use, help is just a few keystrokes away. All UNIX systems have an online manual called **`man`**. Just type **`man {switches} command`**, where **`command`** is the name of the UNIX command you need help on, and the manual pages for that command will be displayed on the screen. The **`man`** command invokes a screen paging utility such as **`more`** to let you scroll forward and back through the text.

Normally, you won't need any switches. However, if you can't remember the name of the command you want you can use the **`-k`** switch to get **`man`** to list the names of commands that relate to a particular task. The format is **`man -k keyword`** where **`keyword`** is a word associated with the function of the command you want. For example, if you couldn't remember the name of the command to change the file access permissions, you could type **`man -k permissions`** and **`man`** will tell you that the command you want is **`chmod`**.

Conclusions

UNIX has a reputation for being arcane, but there's a lot of similarity between it and DOS. Once you've learnt the names of the UNIX commands, and got used to the different syntax, you'll find it easy to switch between the two environments.

Of course, UNIX has many more commands and utilities than it is possible to describe in one article. Once you've investigated further and discovered some of them, you may begin to understand why many programmers who have to work under DOS use products like the MKS Toolkit to give them a UNIX-like environment on their PCs. But even if you don't, you'll probably find an elementary knowledge of UNIX to be a useful skill to have in the years to come.

PCSA

The Author

Julian Moss is a PC consultant and writer. He can be contacted as pcsa@oakworth.demon.co.uk.

Hire, Buy Or Lease: How To Decide

With technology changing so fast, and with today's software requiring more and more powerful machines, is outright purchase of PCs always the best way to go? Tim Phillips offers some advice.

There are three ways to get hold of PCs. You can buy them, you can rent them, or you can lease them. Buying accounts for almost all PC use, a trend that is accelerating with the increasing commoditisation of PC hardware. There will always be exceptions to the rule though.

If you don't have the money to buy machines outright, then it ceases to be an option. It still remains the most cost-effective way to acquire equipment if you have a long-term need: after all, the leasing or hiring companies had to buy the equipment first, and they're not in the charity business. However, the support and maintenance that is often included in the price could be an advantage.

Buying

Buying has advantages that even the most ingenious rental scheme cannot match: the equipment becomes yours, and you can do with it what you like, including upgrades. You can move it where you like too - KPMG estimates that half of all LAN users move office every year.

Once you have the equipment it has value as a capital asset, which will please the bank. The asset value of PCs, though, is small, and they are considered of very low value indeed by their fourth and fifth years.

The biggest advantage when buying is the flexibility it can offer you. Leasing companies typically keep only a small set of models from manufacturers like IBM and Compaq. If you want a cheap clone, or just want an unusual configuration, then buying is the only way to do it in most cases. Also if you need some specialised peripheral or feature, it may simply not be available.

You also have the chance to buy from a reseller with whom you can build a relationship. A leasing company will often - though not always - see its job as complete when the equipment is delivered and installed, and will not be able to give the all-round advice and consultancy that your hardware dealer can give. In the long term, this is at least as important as the capital value of the goods provided.

Support and installation flexibility is the other major advantage of buying. Why take an external support package from the lessor (the person who does the leasing) if you (technically known as the lessee) have the facilities to support in-house? A leasing company may not be overjoyed at the prospect of your technical team reconfiguring its PCs, but if you own them you can do what you like to your equipment.

> *"Rental is designed to last a week or a month. In six months, you will have paid much more than the PC cost to buy, so any job lasting longer than two months starts to make rental look expensive."*

Disadvantages

The disadvantages of buying are obvious: the most important is that you need a lot of cash immediately, or within 30 days. It may be possible to arrange some extended credit with a reseller, but profit margins on PCs are so low for the dealer that few are prepared to offer this type of service. Credit control is worth a lot of money to dealers, so don't be surprised if they are less than enthusiastic about financing your purchase.

The risk of committing large capital sums to new equipment may be hard to justify to either a superior or a bank. PCs don't make money by themselves: they simply put you in a better position to make money. This is rarely a strong argument with which to try to extract funding, as the long-term benefits of buying PCs are impossible to quantify effectively.

The second-hand market for computers is not strong, so when buying assume that you will have the computers for their effective working life. Even if this is only one year, you are likely to have paid out less than for any alternative method of obtaining the equipment, but if it is surplus, rarely used or no longer powerful enough, your obsolete equipment has a small resale value.

The twists and turns of PC buying

> *"Be careful when you rent for indeterminate lengths of time, such as the duration of a development project. A job that was designed to last one month can easily end up lasting longer and, by then, much of the cost benefit will have been lost."*

make any time the wrong time to buy. If you are waiting for a technology issue to resolve, or do not want to make a strategic choice, then buying might give you a headache.

Leasing companies like to compare the tax breaks you can obtain favourably with the tax position when buying computers. This is impossible to generalise as tax laws are country-specific, but the general argument stays the same. You account for your capital assets by depreciating them over time - usually between three and five years. Each year a certain amount of the purchase price can be written off against tax. In subsequent years you have the advantage that PCs you may well have thrown away can continue to be depreciated against your tax bill, but in the short term only a small proportion of the capital cost can be claimed against tax.

Renting

Renting is a great idea for a short-term need. A few years ago the rental market for computers was very small. Now it is booming to match the short-term-needs of computer users. Rental has much of the flexibility of buying without the commitment. Its most common use is for customers who have a short-term upsurge in computer demand but can envisage a time when demand will decrease to current levels again. The cost of buying a computer is too great, and the size of the sum that you need to commit when buying inevitably means that money will be slow to sign off. So, you rent.

Rental agreements are often on-going - you simply hand back the equipment when you finish the job. They are often quick to arrange, too, especially if you are prepared to pay a premium for an emergency delivery. This can be a huge asset in a disaster recovery situation.

More low-key renting can take place when you simply have an excess of work to be done. It is no accident that much of the rental market rents Macintoshes, used for design and publishing work. Just as you could get a freelance designer or layout artist to rush through a job in a couple of days, you can acquire the equipment to do the job at short notice by renting. When the job is finished, there's no need to keep the equipment, and back it goes.

Companies that have periodic bursts of intense activity have found that renting in this way is the most cost-effective (though not necessarily cheap) way to manage computer resources.

Leading Edge

Another advantage to renting is that you can keep a certain proportion of your resources at the leading edge of technology. Rather than have last year's model crunching your numbers,

Who Suits What?

Not everyone can be categorised easily, but there are some categories where renting or leasing are clearly economically justified.

RENT IF...

You need equipment for three months or less.
You need a specialist piece of equipment for a finite period.
You want the maximum tax breaks for a short period.
You need equipment to train staff.
You are putting on an exhibition, presentation or conference and you need extra resources quickly.
Your computers have broken down or been stolen and you need short-term replacements.

LEASE IF...

You have a long-term need but less short-term cash.
You need to drastically expand your IT base quickly.
You value cash in the bank and tax breaks above assets.

BUY IF...

You have the cash or someone will lend it to you cheaply.
You are offered a god deal for buying in bulk.
You are keeping the equipment for longer than three months but less than three years.
You want a cheap brand.
You are not with an established company.
You want consultancy as part of the package.

Hire, Buy or Lease

> *"For any department that needs to be quick on its feet, renting is attractive. It is flexible, fast and, if the costs are high, at least they don't need to be signed off by the boss."*

hire this year's for the duration of the job, and next time you need the best resources, hire something even better. If you want more memory, a larger hard disk or a co-processor, add that to your order. If you find your system is underpowered, call the rental company, who will deliver an upgrade or a new system. Trial and error will never be cheap, but it is cheaper this way.

Often, rental costs are 100% tax deductible, which makes rental a cheaper option than it seems. The main problem with rent though is that it is expensive. You didn't think you got all this without a penalty, did you?

Rental is designed to last a week or a month. In six months you will have paid much more than the PC cost to buy, so any job lasting longer than two months makes rental convenient but expensive. There are some discounts for longer rents, although by the time you have deducted the discount, your rental period will still be expensive.

Be careful when you rent for indeterminate lengths of time, say for example the length of a development project. A job that was designed to last one month can easily end up lasting two, and by then much of the cost benefit will have been spent in rent. Like forgetting to return your rented video tape to the shop, the costs can make renting a poor option compared to gritting your teeth and buying for a limited period.

For fixed short rents, the costs are much more controllable, and renting is a popular way for exhibition organisers and conference administrators, as well as trainers, to lay on large numbers of workstations for two or three days. It is also useful in a non-technical environment, like a hotel, where a delivery and installation service is a bonus, and for in-house training, where dedicated PCs are needed. In both cases, renting gives you an immediate custom-designed network without having to take PCs from existing situations and reconfigure them.

Again, there is a downside to the ability to tap into qualified support people in this way. Often, for rental, delivery and installation are charged separately, and can add significantly to the cost of hire. If you have sufficient resources in-house, doing your own installation may save a considerable sum. However, it's often a waste of time for a support department to spend time maintaining machines that will be on their way back to the rental company after a few weeks.

For any department that needs to be quick on its feet, renting is attractive. It is flexible, fast and if the costs are high, at least they don't need to be signed off by the boss.

Leasing

Leasing is in decline in the PC sector, but looks to be buoyed by some inventive ways of keeping the market going. In some ways it has the disadvantages of renting added to the disadvantages of buying, but it can

True Costs

The cost per seat of a PC is not just the hardware and software cost. Renting or leasing may seem expensive compared to buying, but don't forget that the capital outlay on your computer system is only a fraction of the cost of ownership. The higher cost of leasing a computer for instance may be compensated by your country's tax breaks for renters and the advantage it gives your financial status by freeing up capital or reducing long-term debt.

In cost of ownership terms, US-based consultancy the Gartner Group produced some figures in 1993 showing that the average company spends $40,124 over five years for each PC it owns. Of this cost, only 15% is the cost of the hardware and software, with twice as much spent on hardware as on software, even upgrading your software annually.

Technical support accounts for the same cost - around $6,000, including the capital cost of training the support staff. Administration accounts for 14%, and the remaining 56% - a massive $22,400 - is spent on what Gartner Group calls "end user operations". This includes learning costs, amateur support, users' own software development, and the cost of users fiddling around with their PCs. As this is not an identifiable cost, it is usually ignored when the cost of ownership figures are drawn up.

When considering the higher cost of ownership inherent in leasing or hiring your hardware, factor in the hidden costs. For instance, paying twice as much for your hardware and software by leasing means an overall cost of ownership per workstation will be greater by less than 25%. If leasing enables you to bring down support costs, training costs and to boost productivity by giving you access to hardware which otherwise you couldn't afford, then you can save on the cost of all other aspects of ownership.

Don't just look at the tax issues and the invoices you receive when buying or hiring new equipment. The human costs are just as high.

> *"The price varies depending on what you are leasing, how long for, and what type of credit risk you are to the leasing company. Few will touch startup businesses, and you may find yourself paying more if you are in a high-risk business regardless of the equipment you need."*

work for you in the right situation. Leasing was a method of selling PCs inherited from the mainframers and minicomputer specialists. However, if the method of doing business has not changed, the equipment certainly has. Where a mainframe might be good for 20 years, PCs are obsolete in two, and almost laughable in five.

Under the terms of a lease, you rent the goods for a fixed term, commonly set to either three or five years. As with rental, the ownership of the goods remains with the company leasing the equipment to you. Unlike rental, you can't just give it back. You have to pay your monthly instalments for the full duration of the lease. Sometimes, on payment of a final fee, the equipment becomes yours.

Why would you want to do this? It's a question many former lessees have asked. If you are looking to acquire high-end equipment with a long useful life - an expensive server perhaps, or some sophisticated EPOS system - then the lack of an initial outlay is a bonus. Leasing is the common way for shops to hire their fittings, and this is an area where computer systems are commonly leased. It is substantially cheaper than renting, as the companies always get their leases paid in full, even if you return the equipment early.

Leasing can be budgeted simply, and as it is an operating expense, a lease is usually fully tax deductible. This makes leasing cheaper than it seems, and is the main reason your leasing company will give you when it claims value.

Support and maintenance contracts can also be built into the lease, and you are guaranteed a replacement if the equipment breaks down, so it is also a peace-of-mind alternative for sites where you are unable to provide support.

Costs are calculated per month as a fraction of the value of the goods being leased. A typical quote would mean paying one tenth of the value of the goods each month for three years - about one third the price of rental, but you must pay the full lease cost before you are released from the contract.

The price varies depending on what you are leasing, how long for, and what type of credit risk you are to the leasing company. Few will touch startup businesses, and you may find yourself paying more if you are in a high-risk business, regardless of the equipment you need.

Leasing is also practised by banks and other financial institutions, and a traditional mainframe supplier will probably operate the same leasing business to supply whole companyfuls of PCs on request.

One way that leasing is evolving is the lease-to-buy. At the end of your lease term, you get to keep the equipment, but while it is being leased it belongs to the leasing company. This means that while you are effectively buying it on hire purchase, it is still tax deductible. Check with your leasing company to find out whether you need to insure leased machines against theft or accidental damage.

The obvious problem with lease-to-buy, as the leasing companies are well aware of, is who wants a three-year-old PC? To be honest, its useful life for the company is over, so charging a little extra and letting you keep the kit afterwards is good business.

Software Licences

On the PC, there is no software-only hire or lease market, for obvious reasons. Almost all PC software is so cheap that it would be uneconomical, and the terms of a software licence may forbid hiring and leasing anyway.

However, providing your country's laws allow it, you can rent or lease your hardware with software preinstalled, which for a DTP application for example is a large advantage, and certainly eases the support burden.

If you are attaching your PCs to a network, or if your hired PC is in general use, great care should be taken to keep track of your software licensing. All hired applications cannot be used after the rental or lease period is over, unless you have paid for a full licence.

PCSA

The Author

Tim Phillips is a writer and consultant on PC management and strategy issues.

Supporting OS/2

Tim Sipples presents a selection of the most frequently asked questions about IBM OS/2.

Are there any DOS programs that won't run under OS/2 2.1?

OS/2 2.1 will run virtually all DOS applications in existence, including notorious ones such as Microsoft Flight Simulator, Wing Commander, Maple, MatLab (Version 3.5k or later), AutoCAD (Release 12 requires maintenance version C2 or later), and others. Those that do not run generally fall into the following categories:

1. Programs that use Virtual Control Program Interface (VCPI) memory extenders or other extenders which require direct access to 80386 control registers. Since such applications are also all but incompatible with Windows, most vendors have updates for DPMI compatibility.

2. Applications which attempt to directly address the physical sectors of an OS/2 managed non-removable hard disk drive. Such programs include UnErase in Norton Utilities. OS/2 2.1 has a built-in UnDelete feature.

3. Timing sensitive DOS applications. Certain DOS programs that generate digitized sound through the PC's internal speaker may have distorted sound. High speed, real time data collection may be compromised. These problems can often be minimized or even eliminated using OS/2 2.1's DOS Settings.

4. Certain DOS programming debuggers. DOS applications running under OS/2 2.1 are not permitted to access debug registers DR0-DR7 from a DOS session. Also, DOS debuggers will not be able to set hardware break points, and all read/write operations to debug registers in virtual 8086 mode will be ignored.

DOS-based disk caching software is not required since OS/2 includes a built-in disk cache. Also, the OS/2 CD-ROM file system's cache may be adjusted; see the online Command Reference for details.

Why should I use the HPFS? What does it offer? Does it work with DOS?

HPFS, the High Performance File System, is an installable file system provided with OS/2 which may optionally be used instead of (or alongside) the standard DOS-style FAT file system. HPFS offers long file names (up to 254 characters including the path), contiguous storage of extended attributes, resistance to file fragmentation, improved media error handling, smaller cluster size, support for larger file storage devices (up to 512 GB), and speedier disk operation, particularly on large hard disks in systems with more than 6 MB of RAM.

HPFS is not case sensitive, although it does preserve case in file names.

HPFS is not currently supported on removable media, although some programs (eg, BACKUP) preserve long file names on such FAT disks. Also, native mode DOS cannot access a HPFS partition without a special utility. However, DOS/Windows sessions running under OS/2 can use all files that conform to the 8.3 naming conventions, even if they are stored on HPFS volumes. (FAT is not required for compatibility with DOS and Windows applications.)

What hardware does OS/2 need? Is a PS/2 required?

You need any ISA (AT bus), EISA, VESA LocalBus, PCI, or Microchannel PC compatible, with at least a 386SX CPU, 6 MB of RAM, a 60 MB or larger hard disk (with 15-37 MB free, depending on which features you wish to install), a supported video adapter (CGA, EGA, VGA, SuperVGA, 8514/A, XGA, or third-party driver) with appropriate display, and a high density 3.5" or 5.25" floppy drive for installation. A mouse or other pointing device is strongly recommended.

Allow extra RAM and hard disk space for OS/2-based networking, Extended Services and/or extra system loads (ie, an extraordinary number of large applications running simultaneously).

When calculating hard disk space requirements, subtract space occupied by files already on the hard disk which are functionally included in OS/2 2.1 and may be deleted, eg, DOS, a 386 memory manager, Windows 3.1, Adobe Type Manager with base typefaces.

The performance of OS/2 itself will not be increased appreciably with the addition of a math co-processor. However, certain applications may benefit.

The Workplace Shell (WPS) will not operate with the Monochrome Display Adapter or the Hercules Monochrome Graphics Adapter. Usually the WPS will fail to work with monochrome EGA. However, some EGA adapters (eg, Paradise Monochrome EGA Card, ATI EGA Wonder) will emulate all colour EGA modes on TTL monochrome monitors and, thus, will work with the WPS.

The following dual monitor configurations are supported:

- VGA (or most SuperVGA) with Monochrome Display Adapter (not Hercules)
- 8514/A with VGA
- XGA with VGA
- XGA with one or more other XGA

A secondary monitor is most often used with code debuggers that explicitly support such configurations. The Workplace Shell desktop is currently shown only on the primary display.

OS/2 specifically supports hard drive adapters which conform to the Western Digital chipset interface standard (nearly all MFM, RLL, IDE and ESDI adapters) and Adaptec, Future Domain, DPT and IBM SCSI adapters. OS/2 2.1 drivers for most SCSI adapters, eg, Always, Trantor, Rancho, Procomp, Corel Systems, BusLogic, Seagate, Mylex, CE Infosys, Ciprico, MediaVision ProAudio Spectrum and others are available directly from the adapter manufacturers or from Shareware and Freeware sources.

Generic INT13 support is provided for all other hard disk adapters. This support embraces such devices as Iomega's Bernoulli and SyQuest's removable media products, but for best results contact Iomega or SyQuest for an OS/2 2.1 driver.

OS/2 driver support is available for most CD-ROM drives from Apple, Hitachi, NEC, Panasonic, Pioneer, Sony, Texel, Toshiba and IBM.

Will OS/2 2.1 work with SuperVGA adapters?

OS/2 2.1 contains built-in 256 colour drivers for most SuperVGA adapters based on the following chipsets:

- ATI 28800
- Cirrus Logic CL-GD5422 and CL-GD5424
- Headland Technologies HT209
- IBM VGA256C
- Trident Microsystems TVGA8900B and TVGA8900C
- Tseng Labs ET4000
- Western Digital/Paradise WD90C11, WD90C30, and WD90C31 (in WD90C30 mode)

Some SuperVGA adapters (notably ATI's Vantage and Ultra lines) are 8514/A hardware compatible and will function in 1024x768 256 colour mode with OS/2's built-in 8514/A driver.

Drivers for other SuperVGA adapters (along with installation instructions), and modified versions of the built-in drivers (such as 16 colour versions or small icon versions) should be available from the adapter vendors directly.

IBM has released a shareware set of drivers for SuperVGA adapters based on S3 chipsets.

Regular Windows 3.1 display drivers may be used for the full screen

Actix Graphics Engine 32	[ACTIXS3.ZIP]
Adaptec Virtual ASPI	[VASPBETA.ZIP]
Altrix S3 911, 924, 80x, 928	[S3ALTRIX.ZIP]
Always IN-2000 SCSI	[IN2KADD.ZIP]
AST 4-Port Communications Adapter	[COM01A.ZIP]
ATI Mach8/Mach32 (Ultra and Vantage)	[ATI32V12.ZIP]
BusLogic SCSI	[BT_OS2.ZIP]
Canon BubbleJet Printers	[BJ6OSV12.ZIP]
Canon CaPSL Printers	[CANON38.ZIP]
Cirrus Logic SuperVGA	[CLOS216.ZIP or CLOS2256.ZIP]
Compaq QVision 1024/E, 1024/I	[QVISION.ZIP]
Creative Labs SoundBlaster, SB Pro, SB Pro 2, SB 16	[SB_OS2.ZIP]
Cyrix Cx486DLC Cache Enabler	[CYRIX.ZIP]
Data Technology (DTC) 2270VL and 2130PCI IDE Adapter	[DTCIDE.ZIP]
Dell S3 86C805 On-Board Video	[S3_DELL.ZIP]
Diamond SpeedStar 24X and Western Digital WD90X31 SuperVGA	[SS24XOS2.ZIP]
Diamond Stealth VRAM (S3 911)	[STLOS21.ZIP or STEALTH.ZIP]
Diamond Viper	[VPROS21.ZIP or VPCIOS2.ZIP (PCI) or VPRVLOS2.ZIP (VESA LocalBus)]
Digiboard Intelligent Adapter	[UOS2144.ZIP]
Generic SuperVGA 800x600 16 Color	[SVGA16.ZIP]
Genoa 7000 Series	[GENOA7K.ZIP]
Gravis Ultrasound	[GUSOS201.ZIP or ULTRA03A.ZIP]
Hercules Dynamite/Tseng ET4000-W32(i)	[DYNAMITE.ZIP]
Hewlett-Packard HIL Mouse	[HPMOUSE.ZIP]
IBM Printers (Older Models)	[IBMPRINT.ZIP]
IBM ThinkPad PCMCIA	[TPPCM112.ZIP]
IBM Windsurfer Communications Adapter Telephony	[BMWAVE.ZIP]
LockDrive (Enables HPFS for Magneto-Optical and Similar)	[LOCKDRV.ZIP]
MediaVision ProAudio Mixer and SCSI Port	[MVOS2.ZIP or MVPRODD.ZIP or TMV1SCSI.ZIP]
MediaVision ThunderBoard (and other SoundBlaster Clones)	[TBOS2.ZIP]
Megahertz PCMCIA Modems	[MHZ_PCM.ZIP]
Mitsumi CD-ROM Drive (Most)	[MITFX.ZIP or MITLU002.ZIP]
NCR 77C22/77C22E SuperVGA	[NCRVID21.ZIP]
NEC Pinwriters	[NECPIN.ZIP]
NEC PowerMate 386/25si (Enables Turbo Mode)	[OS2SPD.ZIP]
Number Nine GXE	[NO9OS2.ZIP]
Oak 067 and 077 SuperVGA	[77OS221.ZIP]
Oak 087 ProStar SuperVGA	[87OS221.ZIP]
Orchid F1280(+), 1280-D, VA	[FOS221.ZIP]
Orchid P9000	[P9KOS2_B.ZIP]
PCMCIA Ethernet (Generic)	[ET124PMC.ZIP]
Phillips CM206 CD-ROM	[CM206CDR.ZIP]
Plus Hardcard IIXL	[HRDCD2XL.ZIP]
Promise Disk Accelerator	[PROMISE.ZIP]
Reveal Sound Galaxy (Multimedia Pro 16)	[SC4OS2DV.ZIP]
RNG-810 Random Number Generator	[RNG_810.ZIP]
Rodent (Improved Mouse Driver)	[RODNT100.ZIP]
Roland MPU-401 MIDI	[MPUDD103.ZIP]
S3 Chipset Video Cards (Most)	[S3.ZIP]
Seagate ST-01/02 SCSI	[ST01 102.ZIP]
SIO (Serial Port)	[SIO126E.ZIP]
Sixgraph Wizard 900VL	[P9OS2140.ZIP]
Sony CDU-31A CD-ROM Drive	[SONY31A.ZIP]
Sony CDU-53x CD-ROM Drive	[CDU535.ZIP]
Sony CDU-8002 or Apple CD-150 CD-ROM	[CDU8002A.ZIP]
Spea Video7 Mirage/Mercury	[V7MIRAGE.ZIP or SPEA V7.ZIP]

Figure 1 - Additional OS/2 device drivers available as shareware and freeware

Supporting OS/2

Win-OS/2 desktop. To install a Windows 3.1 display driver under Win-OS/2, simply replace the \OS2\MDOS\WINOS2\SYSTEM\VGA.DRV file with the Windows .DRV file supplied by the vendor. Note that you may have to use the EXPAND program supplied with Win-OS/2 to decompress the vendor's .DRV file.

OS/2 2.1 has standardised the procedure for setting customized refresh rates using DOS-based utilities (such as VMODE). You will be prompted during installation of a SuperVGA driver, and you can specify the DOS utility you wish to execute along with its parameters (if any). You may also need to invoke the utility in your OS/2 AUTOEXEC.BAT file so that DOS sessions are properly configured to your monitor's specifications. Read the OS/2 manuals carefully for details on how to use the DSPINSTL program to install SuperVGA drivers and customize refresh rates for your monitor. Note that for SuperVGA adapters based on the Tseng 4000 chipset only, you may use the command SVGA MONITOR from any full screen DOS session to customize refresh rates for your particular monitor.

Which printers does OS/2 support?

OS/2 2.1 includes support for Hewlett-Packard LaserJets (including the LaserJet 4 series), DeskJets (including the 1200C and 510 models), and PaintJets; IBM/Lexmark ExecJets, InkJets, Proprinters, Quickwriters, Quietwriters, Pageprinters and laser printers; Epson dot matrix, ink jet, and laser printers; PostScript devices; and other printers (eg, Panasonic, Okidata) compatible with these families.

A variety of IBM and HP plotters (including HPGL/2 plotters) is also supported. Drivers for NEC dot matrix printers and Canon laser and BubbleJet printers are also available.

Many drivers can be found on IBM's OS/2 support BBS in the US, on +1 714 438 3325.

Most BubbleJet printers also work quite well with either the IBM 4072 ExecJet or IBM 4070 InkJet driver. If your printer model is not listed when you install OS/2 2.1, consult your Using the Operating System manual starting on Page 351 for advice. Or check your printer's manual to see if the printer emulates one of the models that is listed.

DOS/Windows printer drivers continue to work for DOS/Windows applications. OS/2 2.1 includes a large assortment of Windows printer drivers for Win-OS/2. If necessary install Windows printer drivers using the Win-OS/2 Control Panel. Also, OS/2 2.1 and OS/2 1.3 printer drivers are interchangeable in many cases.

The Lexmark BBS, for IBM printer drivers, is on +1 606 232 5653 in the US.

If you are using a PostScript printer, and you are having difficulty printing under Win-OS/2 3.1, particularly over a network, try adding a "CtrlD=0" line (without the quotes) to the PostScript section of your WIN.INI file.

Can we use COM3 and COM4 in OS/2?

COM3 and COM4 are supported on most PS/2s without any additional effort. On ISA, EISA, PCI, and VESA LocalBus machines, some additions are required to CONFIG.SYS. Using a text editor, include "(port number, base address, interrupt number)" parameters next to the COM.SYS and VCOM.SYS filenames. One example:

```
DEVICE=C:\OS2\COM.SYS
(3,3E8,5) (4,2E8,10)
DEVICE=C:\OS2\MDOS\VCOM.SYS
(3,3E8,5) (4,2E8,10)
```

Parameters for COM1 and COM2 need not be included (unless they are somehow nonstandard). OS/2 must end up with sequentially numbered logical COM ports, if possible. For example, if (physical) port two is not installed but port three or port four is installed, start numbering using (2,...) in the DEVICE lines.

Can OS/2 use more than 16 MB of RAM?

OS/2 2.1 will address all the RAM in your system. If the BIOS recognizes the memory, OS/2 will find it and use it.

However, on certain systems the RAM beyond the 16 MB boundary may be used as a fast swap area. OS/2 relies on its swap file, SWAPPER.DAT, to hold code and data which cannot fit into real memory (ie, to provide virtual memory). If the swap file (and applications) can only be accessed via a hard disk adapter which uses 24-bit DMA for disk access (eg, the Adaptec 154x series), then the system must move code and data below the 16 MB boundary before it can write it to disk. This "double move" is costly (in terms of performance), and often OS/2 will merely use all the RAM above the 16 MB boundary as a fast swap area (before writing to disk) to avoid the problem. It is up to the hard disk adapter driver, however, to decide how to handle this situation.

Only AT bus adapters are limited to 24-bit DMA. Microchannel, EISA, VESA LocalBus, and other 32-bit adapters are not so limited. Moreover, only a select few AT bus hard disk adapters utilize DMA. Nearly all MFM, RLL, ESDI and IDE adapters, and many SCSI adapters, do not use DMA for disk access.

Spider BlackWidow VLB	[BWOS2_2M.ZIP]
Standard PC Speaker (MMPM/2)	[SPEAKER.ZIP]
Summa MM1201 (and compatible) Graphics Tablet	[SUMMA.ZIP]
Tekram DC-620B, DC-660, DC-680 IDE Adapter	[DC620OS2.ZIP]
Toshiba T-4500 PCMCIA Socket Services	[OS2T45.ZIP]
Trantor SCSI	[OS2_TSL6.ZIP]
Trident SuperVGA	[TRID32.ZIP or TRIDENT.ZIP]
Tseng ET-4000/W32, W32i, W32p SuperVGA (Generic)	[OS2W32.ZIP]
Ultrastor 14F	[14OS22X.ZIP]
Ultrastor 24F	[24F_OS2.ZIP and/or 24OS220.ZIP]
Ultrastor 34F	[34F_OS2.ZIP]

Figure 1 - Additional OS/2 device drivers available as shareware and freeware (Continued)

4OS2 Version 1.11 (4OS232.ZIP)
A replacement command interpreter from JP Software. A must for command line users.
BlackHole (BLCKH3.ZIP)
A Workplace Shell object that destroys anything dragged to it.
BookShelf (BOOKSHLF.ZIP)
Utility which presents a coherent menu of all available INF files on your system.
BootOS2 (BOOT2X.ZIP)
Creates an OS/2 bootable diskette.
C-Kermit 5A(189) (CK5A189.ZIP)
A terminal emulation and modem communications program featuring the Kermit file transfer protocol.
CONFIG.SYS Editor (CFGED1B.ZIP)
A Presentation Manager utility which eases CONFIG.SYS editing.
emTeX (various)
Provides professional typesetting and document composition.
EMX/GCC (various)
Powerful C/C++ compiler with programming aids and enhanced libraries.
Enhanced Editor Toolkit and Accessories (various)
Add-ons to the Enhanced Editor (EPM) which provide editor macro capabilities, documentation, and various accessories.
Extended Attributes Backup (EABK202.ZIP)
Saves extended attributes so that non-EA aware backup software preserves all necessary OS/2 data.
Galleria (GALLER.ZIP)
Graphics manipulation and screen capture utility. Also try PM Camera for screen capture.
GhostScript PM (GS252PM.ZIP)
Postscript interpreter and viewer.
GTAK GNU tar (GTAK212B.ZIP)
Tape archive (backup and restore) utility. Supports SCSI tape drives.
HSwitch (HSWTCH02.ZIP)
A task list for full screen sessions.
Icon Extractor (ICON_170.ZIP)
Converts Windows icons to OS/2 format. Icons can be extracted from Windows executables. Assigns icons via drag and drop. Deletes undeleteable objects.
Icon Programming Language (ICON88.ZIP)
A simple yet powerful programming language for many platforms, including OS/2. Supports graphical applications.
Info-Zip's UnZip 5.1 (UNZ51X32.EXE)
Extract files from ZIP archives. PKZip 2.x compatible. Supports extended attributes. Companion utility, Zip 1.9, also available.
INI Maintenance (INIMT21B.ZIP)
Edit and maintain your vital OS/2 INI files.
McAfee's Virus Scan (OSCN113.ZIP)
Detects viruses. Companion Virus Clean and Net Scan utilities also available.
Mousey (MOUSEY10.ZIP)
Shareware utility which allows you to change the default mouse pointer to a candy cane, chicken head, or other icon of your choosing.
Mr. File/PM (MRFILEPM.ZIP)
File manager and program launcher.
OS2Exec (OS2EXEC.ZIP)
Start any OS/2 program from any OS/2 DOS session.
OS2You (OS2YOU27.ZIP)
OS/2 remote control over a modem or LAN connection. Companion program PM2You, for control of graphical applications, including DOS and Windows, is also available.
PMComm (PMCOM110.ZIP)
As full featured as TE/2, but with a Presentation Manager interface.
PMView (PMVU86B.ZIP)
Shareware GIF, JPEG, etc. image viewer with slideshow and drag-and-drop features. Other image viewers include JoeView, Image Archiver, and PMJPEG.

Figure 2 - Popular OS/2 shareware and freeware

If you are planning new hardware purchases, you may wish to take this particular hardware design limitation into account. Specifically, if you plan to install more than 16 MB of RAM in a system, either choose a 32-bit hard disk adapter (Microchannel or EISA, for example) or choose an AT bus adapter which does not utilize DMA for disk access (a standard IDE adapter, an Adaptec 152x series SCSI adapter, or a Future Domain SCSI adapter, for example). The performance trade-off is highly system dependent, however. You may find that even DMA adapters such as the Adaptec 154x series outperform the alternatives in certain cases.

What device drivers are available that aren't included with OS/2's standard package?

Figure 1 shows a list of device drivers (and their common filenames) available from shareware and freeware sources. New and updated OS/2 device drivers are being released almost every day, so use this list only as a rough guide. If you do not see the driver you need listed, you may be able to use one of OS/2's built-in drivers or a "generic" driver (such as the IBM-produced Mitsumi CD-ROM or S3 video card drivers). NDIS drivers for network cards are not included in this list.

Where can I obtain OS/2 shareware and freeware?

On the Internet, the Usenet conference comp.binaries.os2 carries OS/2 software. And several sites are available via anonymous ftp. Some are:

ftp-os2.cdrom.com
ftp-os2.nmsu.edu
software.watson.ibm.com
mtsg.ubc.ca
access.usask.ca
luga.latrobe.edu.au
funic.funet.fi
pdsoft.lancs.ac.uk
ftp.uni-stuttgart.de
src.doc.ic.ac.uk
zaphod.cs.uwindsor.ca
ftp.luth.se
ftp.informatik.tu-muenchen.de

Supporting OS/2

The ftp-os2.cdrom.com library is available on CD-ROM from Walnut Creek (phone +1 510 947 5996). A quarterly subscription is available. Profit Press (+1 800 843 7990) also offers OS/2 shareware and freeware on CD-ROM. EMS (+1 301 924 3594) offers an OS/2 shareware/freeware library on diskette.

What are the "Must Have" Shareware and Freeware titles for OS/2?

Figure 2 lists some of the shareware and freeware selections that have proven popular among OS/2 users. Where available, an approximate filename is provided. However, since version numbers are changing frequently, bear in mind that some of this information may be dated.

What do I need for OS/2 multimedia applications?

OS/2 2.1 includes both the Win-OS/2 multimedia extensions and MMPM/2, the OS/2 multimedia extensions. OS/2 2.1's MMPM/2 includes software motion video support for both IBM Ultimotion and Intel Indeo files. Software motion video provides playback of video clips in a window under OS/2. (Video for Windows will operate correctly under Win-OS/2, but IBM recommends Ultimotion as being more capable than Video for Windows. Ultimotion supports higher frame rates, larger image sizes, better synchronization of video and audio, and, often, simultaneous playback of two or more video clips, even with background tasks running.) An accelerated display adapter and a fast processor can help improve the quality of software motion video.

Drivers for the Creative Labs SoundBlaster series, MediaVision ProAudio series, and IBM M-Audio Capture and Playback Adapter are provided with MMPM/2. Other drivers, including one for the Gravis Ultrasound card, should be available from Shareware and Freeware sources.

A driver for the PC speaker has been released by IBM, but this demands a huge amount of processor attention and does not provide the fidelity that audio adapters do.

Note that the MediaVision ProAudio Spectrum Plus operates correctly when using the built-in MMPM/2 ProAudio Spectrum 16 driver. However, to obtain full functionality you must change the PARAMSTRING line in the [ibmwavepas1601] section of the file \MMOS2\MMPM2.INI after installation of MMPM/2. Using a text editor (like the OS/2 System Editor), change the line so that BPS=8 instead of BPS=16. This change causes MMPM/2 to default to 8-bit audio (since the ProAudio Spectrum Plus does not support 16-bit audio).

The MMPM/2 driver for the original SoundBlaster supplied with OS/2 2.1 contains a special check which may render it inoperable with SoundBlaster clones (like the MediaVision Thunderboard). You can patch the driver using the following commands:

```
CD \MMOS2
COPY SBD2.SYS SBD2.OLD
DEBUG SBD2.SYS
-e 3fd8 00
-w
-q
```

To install a Windows sound driver under Win-OS/2, make sure that the line drivers=mmsystem.dll appears in the file \OS2\MDOS\WINOS2\SYSTEM.INI under the [boot] section. In some cases an installation bug prevents this line from being set up properly.

Should I worry about viruses when running OS/2 2.1?

At present there are only one or two OS/2-specific viruses. However, DOS/Windows viruses can conceivably infect an OS/2 2.1 system. DOS/Windows antivirus tools are just as useful in preventing such infection. Also, IBM has an antivirus package (AntiVirus/2) which runs under OS/2 directly (without DOS/Windows emulation).

What networking products are available for OS/2 2.1?

In addition to DOS/Windows products, OS/2-specific TCP/IP support is available from IBM, Essex Systems (phone +1 508 750 6200), FTP Software (send mail to info@ftp.com), Wollongong (phone +1 415 962 7100), Novell, Ipswitch Inc (+1 617 942 0621) and others. A freeware SLIP implementation, PMNOS, is also available.

Optional add-ons to IBM's TCP/IP package provide X-Window System server and/or client support. Hummingbird Communications offers the eXceed X-Window System server for OS/2.

The NetWare Workstation Kit 2.01 for OS/2 is available from Novell and IBM for a small charge. It is also available free of charge from Shareware and Freeware sources. Note that NetWare 4.01 Server (including all the NLMs) can run alongside OS/2 2.1 at the same time.

The package to enable NetWare 4.01 Server to run alongside OS/2 2.1

SIO COM Drivers (SIO126.ZIP)
Replacement serial port drivers which offer enhanced performance.
StartD (STARTD22.ZIP)
Provides the capability to start DOS sessions with specific, custom DOS Settings from the OS/2 command line.
TE/2 (TE2_124.ZIP)
A full featured terminal emulation and modem communications program.
TinyEd (TINYED.ZIP)
Text editor, under 10K in size. Ideal for an emergency boot diskette.
UUPC/Extended (various)
Provides uucp connection for mail, news, and other services.
Visual REXX (VREXX2.ZIP)
Provides the ability to write REXX programs which use Presentation Manager windows, scroll bars, menus, and other features.
Workplace Shell Backup (WPSBK202.ZIP)
Back up the OS/2 desktop.
Workplace Shell Tools (WPTOOL09.ZIP)
Creates or deletes standard Workplace Shell objects.

Figure 2 - Popular OS/2 shareware and freeware (Continued)

is available from IBM Direct, as part number 53G5730. IBM also offers both NetWare and LAN Server 3.01 (Basic and Advanced) with LAN requesters. Microsoft sells LAN Manager (which comes bundled with OS/2 1.3). An OS/2 2.1 LAN Manager Requester, part no. 96F8359, is available from IBM; Microsoft offers a similar requester with LAN Manager 2.1a (or later). DEC sells Pathworks (Version 2.0b or later) for OS/2. Banyan Vines 5.5 (or later) includes an OS/2 2.x requester.

OS/2 NDIS MAC drivers for most network adapters are available from Shareware and Freeware sources, as well as vax.ftp.com via Internet anonymous ftp. Drivers for 3Com adapters are available via anonymous ftp from ftp.3com.com, from the 3Com Support BBS, or from CompuServe. Drivers for Cabletron adapters are available via anonymous ftp from ftp.ctron.com.

Peer-to-peer networking under OS/2 2.1 is best accomplished with either a TCP/IP or APPN product, like IBM's TCP/IP or Communications Manager/2, respectively. TCP/IP (with NFS) is available for nearly all platforms, so it is a particularly good choice in a mixed environment. Artisoft is reportedly working on an OS/2 version of LANtastic.

Note that LANtastic 5.0 (or above) can run in an OS/2 2.1 DOS session, providing either client or server capability (including access to HPFS long filenames for DOS clients). IBM may have its own offering in the near future (a repackaging and upgrade of the peer-to-peer support already provided in LAN Server 3.01). This support can already interoperate with Microsoft Windows for Workgroups.

What is Extended Services?

Prior to Version 2.0, IBM offered two separate packages with each release of OS/2: Standard Edition and Extended Edition. Extended Edition included extra, bundled software products: the Communications Manager (for communication with IBM mainframes, minicomputers, and other hosts), Database Manager (a network aware, relational database), and LAN Requester.

IBM has now unbundled the Extended Edition features, dropped LAN Requester from the package (now available separately, with IBM's LAN Server), updated it for OS/2 2.x, and renamed it Extended Services 1.0. ES by itself no longer includes the base operating system as Extended Edition once did.

This new arrangement makes it easier to update the base operating system with CSDs (Corrective Service Diskettes). And now ES 1.0 will run under OS/2 1.3 Standard Edition as well as OS/2 2.x. Also, ES 1.0, like OS/2 2.x itself, is designed to operate on both IBM and non-IBM systems.

While Extended Services will remain available, IBM has decided to further separate the ES components. DB2/2, a full 32-bit relational database, is available separately as a replacement for ES's Database Manager. Communications Manager/2 is now also available by itself.

To install Extended Services 1.0 under OS/2 2.1, you must replace the ESSTART.CMD file on ES Diskette 1 with the ESSTART.BAK file found in the \OS2\INSTALL directory. Keep a backup of the old file.

What is the best way to partition a hard disk for OS/2?

It can be useful to place OS/2 by itself in a separate partition of about 40 or 50 MB. If you then elect to experiment with beta releases of OS/2, you can reformat that particular partition quite easily to erase all old code. The swap file can be placed on another partition; its location is determined by the SWAPPATH line in CONFIG.SYS.

Partitioning should be performed by booting the OS/2 Installation Disk. Executing FDISK from the command line will not allow certain functions to be performed.

Sometimes OS/2 2.1 will freeze when I run an application. What do I do?

Before rebooting with CTRL-ALT-DEL, try CTRL-ESC. Do not hit additional keys, do not move the mouse. Wait up to a minute. Either the Window List or an error message should pop up. You may close the offending application at that point; allow some time for it to close. (Try ALT-ESC if you have disabled CTRL-ESC in that application's DOS Settings. If you do not get any response, press CTRL-ESC or ALT-ESC repeatedly until the dialog appears.) Note that the Workplace Shell can recycle, independent of running applications, if it crashes.

If the system is badly disabled, sometimes pressing CTRL-ALT-NUMLOCK twice will result in a prompt to create a dump diskette. You may do so (the online Command Reference describes dump diskettes), or at that point you may reboot with CTRL-ALT-DEL.

To restore the desktop to "factory defaults," use ALT-F1 when OS/2 2.1 starts. See Appendix C of the OS/2 2.1 Using the Operating System manual for details. Note that if you have installed an OS/2 Service Pak the ALT-F1 sequence will restore your desktop to the state it was found in just before installation of the Service Pak.

Acknowledgement

This article is based on the OS/2 Frequently Asked Questions list, compiled and maintained by Timothy F. Sipples, and is used with permission. The author can be contacted via email as usib58c5@ibmmail.com.